arXiv:1411.0900v1 [stat.ML] 4 Nov 2014

Kernel Mean Estimation via Spectral Filtering

Krikamol Muandet
MPI-IS, Tübingen
krikamol@tue.mpg.de

Bharath Sriperumbudur
Dept. of Statistics, PSU
bks18@psu.edu

Bernhard Schölkopf
MPI-IS, Tübingen
bs@tue.mpg.de

Abstract

The problem of estimating the kernel mean in a reproducing kernel Hilbert space (RKHS) is central to kernel methods in that it is used by classical approaches (e.g., when centering a kernel PCA matrix), and it also forms the core inference step of modern kernel methods (e.g., kernel-based non-parametric tests) that rely on embedding probability distributions in RKHSs. Previous work [1] has shown that shrinkage can help in constructing "better" estimators of the kernel mean than the empirical estimator. The present paper studies the consistency and admissibility of the estimators in [1], and proposes a wider class of shrinkage estimators that improve upon the empirical estimator by considering appropriate basis functions. Using the kernel PCA basis, we show that some of these estimators can be constructed using spectral filtering algorithms which are shown to be consistent under some technical assumptions. Our theoretical analysis also reveals a fundamental connection to the kernel-based supervised learning framework. The proposed estimators are simple to implement and perform well in practice.

1 Introduction

The kernel mean or the mean element, which corresponds to the mean of the kernel function in a reproducing kernel Hilbert space (RKHS) computed w.r.t. some distribution $\mathbb{P}$, has played a fundamental role as a basic building block of many kernel-based learning algorithms [2–4], and has recently gained increasing attention through the notion of embedding distributions in an RKHS [5–13]. Estimating the kernel mean remains an important problem as the underlying distribution $\mathbb{P}$ is usually unknown and we must rely entirely on the sample drawn according to $\mathbb{P}$.

Given a random sample drawn independently and identically (i.i.d.) from $\mathbb{P}$, the most common way to estimate the kernel mean is by replacing $\mathbb{P}$ by the empirical measure, $\mathbb{P}_n := \frac{1}{n}\sum_{i=1}^{n} \delta_{X_i}$ where δ_x is a Dirac measure at x [5, 6]. Without any prior knowledge about $\mathbb{P}$, the empirical estimator is possibly the best one can do. However, [1] showed that this estimator can be "improved" by constructing a shrinkage estimator which is a combination of a model with low bias and high variance, and a model with high bias but low variance. Interestingly, significant improvement is in fact possible if the trade-off between these two models is chosen appropriately. The shrinkage estimator proposed in [1], which is motivated from the classical James-Stein shrinkage estimator [14] for the estimation of the mean of a normal distribution, is shown to have a smaller mean-squared error than that of the empirical estimator. These findings provide some support for the conceptual premise that we might be somewhat pessimistic in using the empirical estimator of the kernel mean and there is abundant room for further progress.

In this work, we adopt a spectral filtering approach to obtain shrinkage estimators of kernel mean that improve on the empirical estimator. The motivation behind our approach stems from the idea presented in [1] where the kernel mean estimation is reformulated as an empirical risk minimization (ERM) problem, with the shrinkage estimator being then obtained through penalized ERM. It is important to note that this motivation differs fundamentally from the typical supervised learning as the goal of regularization here is to get the James-Stein-like shrinkage estimators [14] rather than

to prevent overfitting. By looking at regularization from a filter function perspective, in this paper, we show that a wide class of shrinkage estimators for kernel mean can be obtained and that these estimators are consistent for an appropriate choice of the regularization/shrinkage parameter.

Unlike in earlier works [15–18] where the spectral filtering approach has been used in supervised learning problems, we here deal with unsupervised setting and only leverage spectral filtering as a way to construct a shrinkage estimator of the kernel mean. One of the advantages of this approach is that it allows us to incorporate meaningful prior knowledge. The resultant estimators are characterized by the filter function, which can be chosen according to the relevant prior knowledge. Moreover, the spectral filtering gives rise to a broader interpretation of shrinkage through, for example, the notion of early stopping and dimension reduction. Our estimators not only outperform the empirical estimator, but are also simple to implement and computationally efficient.

The paper is organized as follows. In Section 2, we introduce the problem of shrinkage estimation and present a new result that theoretically justifies the shrinkage estimator over the empirical estimator for kernel mean, which improves on the work of [1] while removing some of its drawbacks. Motivated by this result, we consider a general class of shrinkage estimators obtained via spectral filtering in Section 3 whose theoretical properties are presented in Section 4. The empirical performance of the proposed estimators are presented in Section 5. The missing proofs of the results are given in the appendix.

2 Kernel mean shrinkage estimator

In this section, we present preliminaries on the problem of shrinkage estimation in the context of estimating the kernel mean [1] and then present a theoretical justification (see Theorem 1) for shrinkage estimators that improves our understanding of the kernel mean estimation problem, while alleviating some of the issues inherent in the estimator proposed in [1].

Preliminaries: Let $\mathcal{H}$ be an RKHS of functions on a separable topological space $\mathcal{X}$. The space $\mathcal{H}$ is endowed with inner product $\langle \cdot, \cdot \rangle$, associated norm $\| \cdot \|$, and reproducing kernel $k : \mathcal{X} \times \mathcal{X} \to \mathbb{R}$, which we assume to be continuous and bounded, i.e., $\kappa := \sup_{x \in \mathcal{X}} \sqrt{k(x,x)} < \infty$. The kernel mean of some unknown distribution $\mathbb{P}$ on $\mathcal{X}$ and its empirical estimate—we refer to this as *kernel mean estimator* (KME)—from i.i.d. sample $x_1, \ldots, x_n$ are given by

$$\mu_{\mathbb{P}} := \int_{\mathcal{X}} k(x, \cdot) \, \mathrm{d}\mathbb{P}(x) \qquad \text{and} \qquad \hat{\mu}_{\mathbb{P}} := \frac{1}{n} \sum_{i=1}^{n} k(x_i, \cdot), \tag{1}$$

respectively. As mentioned before, $\hat{\mu}_{\mathbb{P}}$ is the "best" possible estimator to estimate $\mu_{\mathbb{P}}$ if nothing is known about $\mathbb{P}$. However, depending on the information that is available about $\mathbb{P}$, one can construct various estimators of $\mu_{\mathbb{P}}$ that perform "better" than $\mu_{\mathbb{P}}$. Usually, the performance measure that is used for comparison is the mean-squared error though alternate measures can be used. Therefore, our main objective is to improve upon KME in terms of the mean-squared error, i.e., construct $\tilde{\mu}_{\mathbb{P}}$ such that $\mathbb{E}_{\mathbb{P}} \|\tilde{\mu}_{\mathbb{P}} - \mu_{\mathbb{P}}\|^2 \leq \mathbb{E}_{\mathbb{P}} \|\hat{\mu}_{\mathbb{P}} - \mu_{\mathbb{P}}\|^2$ for all $\mathbb{P} \in \mathcal{P}$ with strict inequality holding for at least one element in $\mathcal{P}$ where $\mathcal{P}$ is a suitably large class of Borel probability measures on $\mathcal{X}$. Such an estimator $\tilde{\mu}_{\mathbb{P}}$ is said to be *admissible* w.r.t $\mathcal{P}$. If $\mathcal{P} = M_+^1(\mathcal{X})$ is the set of all Borel probability measures on $\mathcal{X}$, then $\tilde{\mu}_{\mathbb{P}}$ satisfying the above conditions may not exist and in that sense, $\hat{\mu}_{\mathbb{P}}$ is possibly the best estimator of $\mu_{\mathbb{P}}$ that one can have.

Admissibility of shrinkage estimator: To improve upon KME, motivated by the James-Stein estimator, $\tilde{\theta}$, [1] proposed a shrinkage estimator $\hat{\mu}_\alpha := \alpha f^* + (1-\alpha)\hat{\mu}_{\mathbb{P}}$ where $\alpha \in \mathbb{R}$ is the shrinkage parameter that balances the low-bias, high-variance model ($\hat{\mu}_{\mathbb{P}}$) with the high-bias, low-variance model ($f^* \in \mathcal{H}$). Assuming for simplicity $f^* = 0$, [1] showed that $\mathbb{E}_{\mathbb{P}} \|\hat{\mu}_\alpha - \mu_{\mathbb{P}}\|^2 < \mathbb{E}_{\mathbb{P}} \|\hat{\mu}_{\mathbb{P}} - \mu_{\mathbb{P}}\|^2$ if and only if $\alpha \in (0, 2\Delta/(\Delta + \|\mu_{\mathbb{P}}\|^2))$ where $\Delta := \mathbb{E}_{\mathbb{P}} \|\hat{\mu}_{\mathbb{P}} - \mu_{\mathbb{P}}\|^2$. While this is an interesting result, the resultant estimator $\hat{\mu}_\alpha$ is strictly not a "statistical estimator" as it depends on quantities that need to be estimated, i.e., it depends on α whose choice requires the knowledge of $\mu_{\mathbb{P}}$, which is the quantity to be estimated. We would like to mention that [1] handles the general case with f^* being not necessarily zero, wherein the range for α then depends on f^* as well. But for the purposes of simplicity and ease of understanding, for the rest of this paper we assume $f^* = 0$. Since $\hat{\mu}_\alpha$ is not practically interesting, [1] resorted to the following representation of $\mu_{\mathbb{P}}$ and $\hat{\mu}_{\mathbb{P}}$ as solutions to the minimization problems [1, 19]:

$$\mu_{\mathbb{P}} = \arg\inf_{g\in\mathcal{H}} \int_{\mathcal{X}} \|k(x,\cdot) - g\|^2 \,\mathrm{d}\mathbb{P}(x), \qquad \hat{\mu}_{\mathbb{P}} = \arg\inf_{g\in\mathcal{H}} \frac{1}{n}\sum_{i=1}^{n} \|k(x_i,\cdot) - g\|^2, \tag{2}$$

using which $\hat{\mu}_\alpha$ is shown to be the solution to the regularized empirical risk minimization problem:

$$\check{\mu}_\lambda = \arg\inf_{g\in\mathcal{H}} \frac{1}{n}\sum_{i=1}^{n} \|k(x_i,\cdot) - g\|^2 + \lambda\|g\|^2, \tag{3}$$

where $\lambda > 0$ and $\alpha := \frac{\lambda}{\lambda+1}$, i.e., $\check{\mu}_\lambda = \hat{\mu}_{\frac{\lambda}{\lambda+1}}$. It is interesting to note that unlike in supervised learning (e.g., least squares regression), the empirical minimization problem in (2) is not ill-posed and therefore does not require a regularization term although it is used in (3) to obtain a shrinkage estimator of $\mu_{\mathbb{P}}$. [1] then obtained a value for λ through cross-validation and used it to construct $\hat{\mu}_{\frac{\lambda}{\lambda+1}}$ as an estimator of $\mu_{\mathbb{P}}$, which is then shown to perform empirically better than $\hat{\mu}_{\mathbb{P}}$. However, no theoretical guarantees including the basic requirement of $\hat{\mu}_{\frac{\lambda}{\lambda+1}}$ being consistent are provided. In fact, because λ is data-dependent, the above mentioned result about the improved performance of $\hat{\mu}_\alpha$ over a range of α does not hold as such a result is proved assuming α is a constant and does not depend on the data. While it is clear that the regularizer in (3) is not needed to make (2) well-posed, the role of λ is not clear from the point of view of $\hat{\mu}_{\frac{\lambda}{\lambda+1}}$ being consistent and better than $\hat{\mu}_{\mathbb{P}}$. The following result provides a theoretical understanding of $\hat{\mu}_{\frac{\lambda}{\lambda+1}}$ from these viewpoints.

Theorem 1. *Let $\check{\mu}_\lambda$ be constructed as in (3). Then the following hold.*

(i) $\|\check{\mu}_\lambda - \mu_{\mathbb{P}}\| \xrightarrow{\mathbb{P}} 0$ as $\lambda \to 0$ and $n \to \infty$. In addition, if $\lambda = n^{-\beta}$ for some $\beta > 0$, then $\|\check{\mu}_\lambda - \mu_{\mathbb{P}}\| = O_{\mathbb{P}}(n^{-\min\{\beta,1/2\}})$.

(ii) For $\lambda = cn^{-\beta}$ with $c > 0$ and $\beta > 1$, define $\mathcal{P}_{c,\beta} := \{\mathbb{P} \in M_+^1(\mathcal{X}) : \|\mu_{\mathbb{P}}\|^2 < A\int k(x,x)\,\mathrm{d}\mathbb{P}(x)\}$ where $A := \frac{2^{1/\beta}\beta}{2^{1/\beta}\beta + c^{1/\beta}(\beta-1)^{(\beta-1)/\beta}}$. Then $\forall\, n$ and $\forall\, \mathbb{P} \in \mathcal{P}_{c,\beta}$, we have $\mathbb{E}_{\mathbb{P}}\|\check{\mu}_\lambda - \mu_{\mathbb{P}}\|^2 < \mathbb{E}_{\mathbb{P}}\|\hat{\mu}_{\mathbb{P}} - \mu_{\mathbb{P}}\|^2$.

Remark. (*i*) Theorem 1(i) shows that $\check{\mu}_\lambda$ is a consistent estimator of $\mu_{\mathbb{P}}$ as long as $\lambda \to 0$ and the convergence rate in probability of $\|\check{\mu}_\lambda - \mu_{\mathbb{P}}\|$ is determined by the rate of convergence of λ to zero, with the best possible convergence rate being $n^{-1/2}$. Therefore to attain a fast rate of convergence, it is instructive to choose λ such that $\lambda\sqrt{n} \to 0$ as $\lambda \to 0$ and $n \to \infty$.

(*ii*) Suppose for some $c > 0$ and $\beta > 1$, we choose $\lambda = cn^{-\beta}$, which means the resultant estimator $\check{\mu}_\lambda$ is a proper estimator as it does not depend on any unknown quantities. Theorem 1(ii) shows that for any n and $\mathbb{P} \in \mathcal{P}_{c,\beta}$, $\check{\mu}_\lambda$ is a "better" estimator than $\hat{\mu}_{\mathbb{P}}$. Note that for any $\mathbb{P} \in M_+^1(\mathcal{X})$, $\|\mu_{\mathbb{P}}\|^2 = \int\int k(x,y)\,\mathrm{d}\mathbb{P}(x)\,\mathrm{d}\mathbb{P}(y) \le (\int \sqrt{k(x,x)}\,\mathrm{d}\mathbb{P}(x))^2 \le \int k(x,x)\,\mathrm{d}\mathbb{P}(x)$. This means $\check{\mu}_\lambda$ is admissible if we restrict $M_+^1(\mathcal{X})$ to $\mathcal{P}_{c,\beta}$ which considers only those distributions for which $\|\mu_{\mathbb{P}}\|^2 / \int k(x,x)\,\mathrm{d}\mathbb{P}(x)$ is strictly less than a constant, $A < 1$. It is obvious to note that if c is very small or β is very large, then A gets closer to one and $\check{\mu}_\lambda$ behaves almost like $\hat{\mu}_{\mathbb{P}}$, thereby matching with our intuition.

(*iii*) A nice interpretation for $\mathcal{P}_{c,\beta}$ can be obtained as in Theorem 1(ii) when k is a translation invariant kernel on $\mathbb{R}^d$. It can be shown that $\mathcal{P}_{c,\beta}$ contains the class of all probability measures whose characteristic function has an L^2 norm (and therefore is the set of square integrable probability densities if $\mathbb{P}$ has a density w.r.t. the Lebesgue measure) bounded by a constant that depends on c, β and k (see §B in the appendix). ■

3 Spectral kernel mean shrinkage estimator

Let us return to the shrinkage estimator $\hat{\mu}_\alpha$ considered in [1], i.e., $\hat{\mu}_\alpha = \alpha f^* + (1-\alpha)\hat{\mu}_{\mathbb{P}} = \alpha\sum_i \langle f^*, e_i\rangle e_i + (1-\alpha)\sum_i \langle \hat{\mu}_{\mathbb{P}}, e_i\rangle e_i$, where $(e_i)_{i\in\mathbb{N}}$ are the countable orthonormal basis (ONB) of $\mathcal{H}$—countable ONB exist since $\mathcal{H}$ is separable which follows from $\mathcal{X}$ being separable and k being continuous [20, Lemma 4.33]. This estimator can be generalized by considering the shrinkage estimator $\hat{\mu}_{\boldsymbol{\alpha}} := \sum_i \alpha_i \langle f^*, e_i\rangle e_i + \sum_i (1-\alpha_i)\langle \hat{\mu}_{\mathbb{P}}, e_i\rangle e_i$ where $\boldsymbol{\alpha} := (\alpha_1, \alpha_2, \ldots) \in \mathbb{R}^\infty$ is a sequence of shrinkage parameters. If $\Delta_{\boldsymbol{\alpha}} := \mathbb{E}_{\mathbb{P}}\|\hat{\mu}_{\boldsymbol{\alpha}} - \mu_{\mathbb{P}}\|^2$ is the risk of this estimator, the following theorem gives an optimality condition on $\boldsymbol{\alpha}$ for which $\Delta_{\boldsymbol{\alpha}} < \Delta$.

Theorem 2. *For some ONB $(e_i)_i$, $\Delta_{\boldsymbol{\alpha}} - \Delta = \sum_i (\Delta_{\boldsymbol{\alpha},i} - \Delta_i)$ where $\Delta_{\boldsymbol{\alpha},i}$ and Δ_i denote the risk of the ith component of $\hat{\mu}_{\boldsymbol{\alpha}}$ and $\hat{\mu}_{\mathbb{P}}$, respectively. Then, $\Delta_{\boldsymbol{\alpha},i} - \Delta_i < 0$ if*

$$0 < \alpha_i < \frac{2\Delta_i}{\Delta_i + (f_i^* - \mu_i)^2}, \tag{4}$$

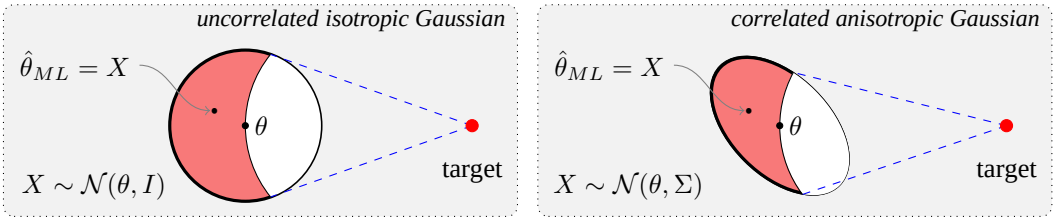

Figure 1: Geometric explanation of a shrinkage estimator when estimating a mean of a Gaussian distribution. For isotropic Gaussian, the level sets of the joint density of $\hat{\theta}_{ML} = X$ are hyperspheres. In this case, shrinkage has the same effect regardless of the direction. Shaded area represents those estimates that get closer to θ after shrinkage. For anisotropic Gaussian, the level sets are concentric ellipsoids, which makes the effect dependent on the direction of shrinkage.

where f_i^* and μ_i denote the Fourier coefficients of f^* and $\mu_{\mathbb{P}}$, respectively.

The condition in (4) is a component-wise version of the condition given in [1, Theorem 1] for a class of estimators $\hat{\mu}_\alpha := \alpha f^* + (1-\alpha)\hat{\mu}_{\mathbb{P}}$ which may be expressed here by assuming that we have a constant shrinkage parameter $\alpha_i = \alpha$ for all i. Clearly, as the optimal range of α_i may vary across coordinates, the class of estimators in [1] does not allow us to adjust α_i accordingly. To understand why this property is important, let us consider the problem of estimating the mean of Gaussian distribution illustrated in Figure 1. For correlated random variable $X \sim \mathcal{N}(\theta, \Sigma)$, a natural choice of basis is the set of orthonormal eigenvectors which diagonalize the covariance matrix Σ of X. Clearly, the optimal range of α_i depends on the corresponding eigenvalues. Allowing for different basis $(e_i)_i$ and shrinkage parameter α_i opens up a wide range of strategies that can be used to construct "better" estimators.

A natural strategy under this representation is as follows: *i*) we specify the ONB $(e_i)_i$ and project $\hat{\mu}_{\mathbb{P}}$ onto this basis. *ii*) we shrink each $\hat{\mu}_i$ independently according to a pre-defined shrinkage rule. *iii*) the shrinkage estimate is reconstructed as a superposition of the resulting components. In other words, an ideal shrinkage estimator can be defined formally as a non-linear mapping:

$$\hat{\mu}_{\mathbb{P}} \longrightarrow \sum_i h(\alpha_i)\langle f^*, e_i\rangle e_i + \sum_i (1 - h(\alpha_i))\langle \hat{\mu}_{\mathbb{P}}, e_i\rangle e_i \tag{5}$$

where $h : \mathbb{R} \to \mathbb{R}$ is a shrinkage rule. Since we make no reference to any particular basis $(e_i)_i$, nor to any particular shrinkage rule h, a wide range of strategies can be adopted here. For example, we can view *whitening* as a special case in which f^* is the data average $\frac{1}{n}\sum_{i=1}^n x_i$ and $1 - h(\alpha_i) = 1/\sqrt{\alpha_i}$ where α_i and e_i are the ith eigenvalue and eigenvector of the covariance matrix, respectively.

Inspired by Theorem 2, we adopt the spectral filtering approach as one of the strategies to construct the estimators of the form (5). To this end, owing to the regularization interpretation in (3), we consider estimators of the form $\sum_{i=1}^n \beta_i k(x_i, \cdot)$ for some $\boldsymbol{\beta} \in \mathbb{R}^n$—looking for such an estimator is equivalent to learning a *signed measure* that is supported on $(x_i)_{i=1}^n$. Since $\sum_{i=1}^n \beta_i k(x_i, \cdot)$ is a minimizer of (3), β should satisfy $\mathbf{K}\boldsymbol{\beta} = \mathbf{K}\mathbf{1}_n$ where $\mathbf{K}$ is an $n \times n$ Gram matrix and $\mathbf{1}_n = [1/n. \ldots, 1/n]^\top$. Here the solution is trivially $\boldsymbol{\beta} = \mathbf{1}_n$, i.e., the coefficients of the standard estimator $\hat{\mu}_{\mathbb{P}}$ if $\mathbf{K}$ is invertible. Since $\mathbf{K}^{-1}$ may not exist and even if it exists, the computation of it can be numerically unstable, the idea of spectral filtering—this is quite popular in the theory of inverse problems [15] and has been used in kernel least squares [17]—is to replace $\mathbf{K}^{-1}$ by some regularized matrices $g_\lambda(\mathbf{K})$ that approximates $\mathbf{K}^{-1}$ as λ goes to zero. Note that unlike in (3), the regularization is quite important here (i.e., the case of estimators of the form $\sum_{i=1}^n \beta_i k(x_i, \cdot)$) without which the the linear system is under determined. Therefore, we propose the following class of estimators:

$$\hat{\mu}_\lambda := \sum_{i=1}^n \beta_i k(x_i, \cdot) \quad \text{with} \quad \boldsymbol{\beta}(\lambda) := g_\lambda(\mathbf{K})\mathbf{K}\mathbf{1}_n, \tag{6}$$

where $g_\lambda(\cdot)$ is a filter function and λ is referred to as a shrinkage parameter. The matrix-valued function $g_\lambda(\mathbf{K})$ can be described by a scalar function $g_\lambda : [0, \kappa^2] \to \mathbb{R}$ on the spectrum of $\mathbf{K}$. That is, if $\mathbf{K} = \mathbf{U}\mathbf{D}\mathbf{U}^\top$ is the eigen-decomposition of $\mathbf{K}$ where $\mathbf{D} = \operatorname{diag}(\tilde{\gamma}_1, \ldots, \tilde{\gamma}_n)$, we have $g_\lambda(\mathbf{D}) = \operatorname{diag}(g_\lambda(\tilde{\gamma}_1), \ldots, g_\lambda(\tilde{\gamma}_n))$ and $g_\lambda(\mathbf{K}) = \mathbf{U}g_\lambda(\mathbf{D})\mathbf{U}^\top$. For example, the scalar filter function of Tikhonov regularization is $g_\lambda(\gamma) = 1/(\gamma + \lambda)$. In the sequel, we call this class of estimators a *spectral kernel mean shrinkage estimator* (Spectral-KMSE).

Table 1: Update equations for $\boldsymbol{\beta}$ and corresponding filter functions.

Algorithm	Update Equation ($\mathbf{a} := \mathbf{K}\mathbf{1}_n - \mathbf{K}\boldsymbol{\beta}^{t-1}$)	Filter Function
L2 Boosting	$\boldsymbol{\beta}^t \leftarrow \boldsymbol{\beta}^{t-1} + \eta\mathbf{a}$	$g(\gamma) = \eta\sum_{i=1}^{t-1}(1-\eta\gamma)^i$
Acc. L2 Boosting	$\boldsymbol{\beta}^t \leftarrow \boldsymbol{\beta}^{t-1} + \omega_t(\boldsymbol{\beta}^{t-1} - \boldsymbol{\beta}^{t-2}) + \frac{\kappa_t}{n}\mathbf{a}$	$g(\gamma) = p_t(\gamma)$
Iterated Tikhonov	$(\mathbf{K} + n\lambda\mathbf{I})\boldsymbol{\beta}_i = \mathbf{1}_n + n\lambda\boldsymbol{\beta}_{i-1}$	$g(\gamma) = \frac{(\gamma+\lambda)^t - \gamma^t}{\lambda(\gamma+\lambda)^t}$
Truncated SVD	None	$g(\gamma) = \gamma^{-1}\mathbb{1}_{\{\gamma \geq \lambda\}}$

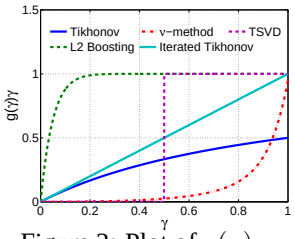

Figure 2: Plot of $g(\gamma)\gamma$.

Proposition 3. *The Spectral-KMSE satisfies* $\hat{\mu}_\lambda = \sum_{i=1}^n g_\lambda(\tilde{\gamma}_i)\tilde{\gamma}_i\langle\hat{\mu}, \tilde{\mathbf{v}}_i\rangle\tilde{\mathbf{v}}_i$*, where* $(\tilde{\gamma}_i, \tilde{\mathbf{v}}_i)$ *are eigenvalue and eigenfunction pairs of the empirical covariance operator* $\widehat{\mathcal{C}}_k : \mathcal{H} \to \mathcal{H}$ *defined as* $\widehat{\mathcal{C}}_k = \frac{1}{n}\sum_{i=1}^n k(\cdot, x_i) \otimes k(\cdot, x_i)$.

By virtue of Proposition 3, if we choose $1 - h(\tilde{\gamma}) := g_\lambda(\tilde{\gamma})\tilde{\gamma}$, the Spectral-KMSE is indeed in the form of (5) when $f^* = 0$ and $(e_i)_i$ is the kernel PCA (KPCA) basis, with the filter function g_λ determining the shrinkage rule. Since by definition $g_\lambda(\tilde{\gamma}_i)$ approaches the function $1/\tilde{\gamma}_i$ as λ goes to 0, the function $g_\lambda(\tilde{\gamma}_i)\tilde{\gamma}_i$ approaches 1 (no shrinkage). As the value of λ increases, we have more shrinkage because the value of $g_\lambda(\tilde{\gamma}_i)\tilde{\gamma}_i$ deviates from 1, and the behavior of this deviation depends on the filter function g_λ. For example, we can see that Proposition 3 generalizes Theorem 2 in [1] where the filter function is $g_\lambda(\mathbf{K}) = (\mathbf{K} + n\lambda\mathbf{I})^{-1}$, i.e., $g(\gamma) = 1/(\gamma + \lambda)$. That is, we have $g_\lambda(\tilde{\gamma}_i)\tilde{\gamma}_i = \tilde{\gamma}_i/(\tilde{\gamma}_i + \lambda)$, implying that the effect of shrinkage is relatively larger in the low-variance direction. In the following, we discuss well-known examples of spectral filtering algorithms obtained by various choices of g_λ. Update equations for $\boldsymbol{\beta}(\lambda)$ and corresponding filter functions are summarized in Table 1. Figure 2 illustrates the behavior of these filter functions.

L2 Boosting. This algorithm, also known as gradient descent or Landweber iteration, finds a weight $\boldsymbol{\beta}$ by performing a gradient descent iteratively. Thus, we can interpret *early stopping* as shrinkage and the reciprocal of iteration number as shrinkage parameter, i.e., $\lambda \approx 1/t$. The step-size η does not play any role for shrinkage [16], so we use the fixed step-size $\eta = 1/\kappa^2$ throughout.

Accelerated L2 Boosting. This algorithm, also known as ν-method, uses an accelerated gradient descent step, which is faster than L2 Boosting because we only need $\sqrt{t}$ iterations to get the same solution as the L2 Boosting would get after t iterations. Consequently, we have $\lambda \approx 1/t^2$.

Iterated Tikhonov. This algorithm can be viewed as a combination of Tikhonov regularization and gradient descent. Both parameters λ and t play the role of shrinkage parameter.

Truncated Singular Value Decomposition. This algorithm can be interpreted as a projection onto the first principal components of the KPCA basis. Hence, we may interpret *dimensionality reduction* as shrinkage and the size of reduced dimension as shrinkage parameter. This approach has been used in [21] to improve the kernel mean estimation under the low-rank assumption.

Most of the above spectral filtering algorithms allow to compute the coefficients $\boldsymbol{\beta}$ without explicitly computing the eigen-decomposition of $\mathbf{K}$, as we can see in Table 1, and some of which may have no natural interpretation in terms of regularized risk minimization. Lastly, an initialization of $\boldsymbol{\beta}$ corresponds to the target of shrinkage. In this work, we assume that $\boldsymbol{\beta}^0 = 0$ throughout.

4 Theoretical properties of Spectral-KMSE

This section presents some theoretical properties for the proposed Spectral-KMSE in (6). To this end, we first present a regularization interpretation that is different from the one in (3) which involves learning a smooth operator from $\mathcal{H}$ to $\mathcal{H}$ [22]. This will be helpful to investigate the consistency of the Spectral-KMSE. Let us consider the following regularized risk minimization problem,

$$\arg\min_{\mathbf{F} \in \mathcal{H}\otimes\mathcal{H}} \quad \mathbb{E}_X \left\| k(X, \cdot) - \mathbf{F}[k(X, \cdot)] \right\|_{\mathcal{H}}^2 + \lambda \|\mathbf{F}\|_{HS}^2 \tag{7}$$

where $\mathbf{F}$ is a Hilbert-Schmidt operator from $\mathcal{H}$ to $\mathcal{H}$. Essentially, we are seeking a smooth operator $\mathbf{F}$ that maps $k(x, \cdot)$ to itself, where (7) is an instance of the regression framework in [22]. The formulation of shrinkage as the solution of a smooth operator regression, and the empirical solution (8) and in the lines below, were given in a personal communication by Arthur Gretton. It can be

shown that the solution to (7) is given by $\mathbf{F} = \mathcal{C}_k(\mathcal{C}_k + \lambda \mathbf{I})^{-1}$ where $\mathcal{C}_k : \mathcal{H} \to \mathcal{H}$ is a covariance operator in $\mathcal{H}$ defined as $\mathcal{C}_k = \int k(\cdot, x) \otimes k(\cdot, x) \, \mathrm{d}\mathbb{P}(x)$ (see §E of the appendix for a proof). Define $\mu_\lambda := \mathbf{F}\mu_{\mathbb{P}} = \mathcal{C}_k(\mathcal{C}_k + \lambda \mathbf{I})^{-1}\mu_{\mathbb{P}}$. Since k is bounded, it is easy to verify that $\mathcal{C}_k$ is Hilbert-Schmidt and therefore compact. Hence by the Hilbert-Schmidt theorem, $\mathcal{C}_k = \sum_i \gamma_i \langle \cdot, \psi_i \rangle \psi_i$ where $(\gamma_i)_{i \in \mathbb{N}}$ are the positive eigenvalues and $(\psi_i)_{i \in \mathbb{N}}$ are the corresponding eigenvectors that form an ONB for the range space of $\mathcal{C}_k$ denoted as $\mathcal{R}(\mathcal{C}_k)$. This implies μ_λ can be decomposed as $\mu_\lambda = \sum_{i=1}^{\infty} \frac{\gamma_i}{\gamma_i + \lambda} \langle \mu_{\mathbb{P}}, \psi_i \rangle \psi_i$. We can observe that the filter function corresponding to the problem (7) is $g_\lambda(\gamma) = 1/(\gamma + \lambda)$. By extending this approach to other filter functions, we obtain $\mu_\lambda = \sum_{i=1}^{\infty} \gamma_i g_\lambda(\gamma_i) \langle \mu_{\mathbb{P}}, \psi_i \rangle \psi_i$ which is equivalent to $\mu_\lambda = \mathcal{C}_k g_\lambda(\mathcal{C}_k)\mu_{\mathbb{P}}$.

Since $\mathcal{C}_k$ is a compact operator, the role of filter function g_λ is to regularize the inverse of $\mathcal{C}_k$. In standard supervised setting, the explicit form of the solution is $f_\lambda = g_\lambda(L_k)L_k f_\rho$ where L_k is the integral operator of kernel k acting in $L^2(\mathcal{X}, \rho_X)$ and f_ρ is the expected solution given by $f_\rho(x) = \int_{\mathcal{Y}} y \, \mathrm{d}\rho(y|x)$ [16]. It is interesting to see that μ_λ admits a similar form to that of f_λ, but it is written in term of covariance operator $\mathcal{C}_k$ instead of the integral operator L_k. Moreover, the solution to (7) is also in a similar form to the regularized conditional embedding $\mu_{Y|X} = \mathcal{C}_{YX}(\mathcal{C}_k + \lambda \mathbf{I})^{-1}$ [9]. This connection implies that the spectral filtering may be applied more broadly to improve the estimation of conditional mean embedding, i.e., $\mu_{Y|X} = \mathcal{C}_{YX} g_\lambda(\mathcal{C}_k)$.

The empirical counterpart of (7) is given by

$$\arg\min_{\mathbf{F}} \quad \frac{1}{n} \sum_{i=1}^{n} \|k(x_i, \cdot) - \mathbf{F}[k(x_i, \cdot)]\|_{\mathcal{H}}^2 + \lambda \|\mathbf{F}\|_{HS}^2, \tag{8}$$

resulting in $\hat{\mu}_\lambda = \mathbf{F}\hat{\mu}_{\mathbb{P}} = \mathbf{1}_n^\top \mathbf{K}(\mathbf{K} + \lambda \mathbf{I})^{-1}\Phi$ where $\Phi = [k(x_1, \cdot), \ldots, k(x_n, \cdot)]^\top$, which matches with the one in (6) with $g_\lambda(\mathbf{K}) = (\mathbf{K} + \lambda \mathbf{I})^{-1}$. Note that this is exactly the F-KMSE proposed in [1]. Based on μ_λ which depends on $\mathbb{P}$, an empirical version of it can be obtained by replacing $\mathcal{C}_k$ and $\mu_{\mathbb{P}}$ with their empirical estimators leading to $\tilde{\mu}_\lambda = \widehat{\mathcal{C}}_k g_\lambda(\widehat{\mathcal{C}}_k)\hat{\mu}_{\mathbb{P}}$. The following result shows that $\hat{\mu}_\lambda = \tilde{\mu}_\lambda$, which means the Spectral-KMSE proposed in (6) is equivalent to solving (8).

Proposition 4. *Let $\widehat{\mathcal{C}}_k$ and $\hat{\mu}_{\mathbb{P}}$ be the sample counterparts of $\mathcal{C}_k$ and $\mu_{\mathbb{P}}$ given by $\widehat{\mathcal{C}}_k := \frac{1}{n}\sum_{i=1}^{n} k(x_i, \cdot) \otimes k(x_i, \cdot)$ and $\hat{\mu}_{\mathbb{P}} := \frac{1}{n}\sum_{i=1}^{n} k(x_i, \cdot)$, respectively. Then, we have that $\tilde{\mu}_\lambda := \widehat{\mathcal{C}}_k g_\lambda(\widehat{\mathcal{C}}_k)\hat{\mu}_{\mathbb{P}} = \hat{\mu}_\lambda$, where $\hat{\mu}_\lambda$ is defined in (6).*

Having established a regularization interpretation for $\hat{\mu}_\lambda$, it is of interest to study the consistency and convergence rate of $\hat{\mu}_\lambda$ similar to KMSE in Theorem 1. Our main goal here is to derive convergence rates for a broad class of algorithms given a set of sufficient conditions on the filter function, g_λ. We believe that for some algorithms it is possible to derive the best achievable bounds, which requires ad-hoc proofs for each algorithm. To this end, we provide a set of conditions any *admissible* filter function, g_λ must satisfy.

Definition 1. *A family of filter functions $g_\lambda : [0, \kappa^2] \to \mathbb{R}, 0 < \lambda \leq \kappa^2$ is said to be admissible if there exists finite positive constants B, C, D, and η_0 (all independent of λ) such that $(C1)$ $\sup_{\gamma \in [0,\kappa^2]} |\gamma g_\lambda(\gamma)| \leq B$, $(C2)$ $\sup_{\gamma \in [0,\kappa^2]} |r_\lambda(\gamma)| \leq C$ and $(C3)$ $\sup_{\gamma \in [0,\kappa^2]} |r_\lambda(\gamma)|\gamma^\eta \leq D\lambda^\eta$, $\forall \, \eta \in (0, \eta_0]$ hold, where $r_\lambda(\gamma) := 1 - \gamma g_\lambda(\gamma)$.*

These conditions are quite standard in the theory of inverse problems [15, 23]. The constant η_0 is called the *qualification* of g_λ and is a crucial factor that determines the rate of convergence in inverse problems. As we will see below, that the rate of convergence of $\hat{\mu}_\lambda$ depends on two factors: (a) smoothness of $\mu_{\mathbb{P}}$ which is usually unknown as it depends on the unknown $\mathbb{P}$ and (b) qualification of g_λ which determines how well the smoothness of $\mu_{\mathbb{P}}$ is captured by the spectral filter, g_λ.

Theorem 5. *Suppose g_λ is admissible in the sense of Definition 1. Let $\kappa = \sup_{x \in \mathcal{X}} \sqrt{k(x, x)}$. If $\mu_{\mathbb{P}} \in \mathcal{R}(\mathcal{C}_k^\beta)$ for some $\beta > 0$, then for any $\delta > 0$, with probability at least $1 - 3e^{-\delta}$,*

$$\|\hat{\mu}_\lambda - \mu_{\mathbb{P}}\| \leq \frac{2\kappa B + \kappa B\sqrt{2\delta}}{\sqrt{n}} + D\lambda^{\min\{\beta, \eta_0\}} \|\mathcal{C}_k^{-\beta}\mu_{\mathbb{P}}\| + C\tau \frac{(2\sqrt{2}\kappa^2\sqrt{\delta})^{\min\{1,\beta\}}}{n^{\min\{1/2, \beta/2\}}} \|\mathcal{C}_k^{-\beta}\mu_{\mathbb{P}}\|,$$

where $\mathcal{R}(A)$ denotes the range space of A and τ is some universal constant that does not depend on λ and n. Therefore, $\|\hat{\mu}_\lambda - \mu_{\mathbb{P}}\| = O_{\mathbb{P}}(n^{-\min\{1/2, \beta/2\}})$ with $\lambda = o(n^{-\frac{\min\{1/2, \beta/2\}}{\min\{\beta, \eta_0\}}})$.

Theorem 5 shows that the convergence rate depends on the smoothness of $\mu_{\mathbb{P}}$ which is imposed through the range space condition that $\mu_{\mathbb{P}} \in \mathcal{R}(\mathcal{C}_k^\beta)$ for some $\beta > 0$. Note that this is in contrast

to the estimator in Theorem 1 which does not require any smoothness assumptions on $\mu_{\mathbb{P}}$. It can be shown that the smoothness of $\mu_{\mathbb{P}}$ increases with increase in β. This means, irrespective of the smoothness of $\mu_{\mathbb{P}}$ for $\beta > 1$, the best possible convergence rate is $n^{-1/2}$ which matches with that of KMSE in Theorem 1. While the qualification η_0 does not seem to directly affect the rates, it controls the rate at which λ converges to zero. For example, if $g_\lambda(\gamma) = 1/(\gamma + \lambda)$ which corresponds to Tikhonov regularization, it can be shown that $\eta_0 = 1$ which means for $\beta > 1$, $\lambda = o(n^{-1/2})$ implying that λ cannot decay to zero slower than $n^{-1/2}$. Ideally, one would require a larger η_0 (preferably infinity which is the case with truncated SVD) so that the convergence of λ to zero can be made arbitrarily slow if β is large. This way, both β and η_0 control the behavior of the estimator.

In fact, Theorem 5 provides a choice for λ—which is what we used in Theorem 1 to study the admissibility of $\check{\mu}_\lambda$ to $\mathcal{P}_{c,\beta}$—to construct the Spectral-KMSE. However, this choice of λ depends on β which is not known in practice (although η_0 is known as it is determined by the choice of g_λ). Therefore, λ is usually learnt from data through cross-validation or through Lepski's method [24] for which guarantees similar to the one presented in Theorem 5 can be provided. However, irrespective of the data-dependent/independent choice for λ, checking for the admissibility of Spectral-KMSE (similar to the one in Theorem 1) is very difficult and we intend to consider it in future work.

5 Empirical studies

Synthetic data. Given the i.i.d. sample $\mathbf{X} = \{x_1, x_2, \ldots, x_n\}$ from $\mathbb{P}$ where $x_i \in \mathbb{R}^d$, we evaluate different estimators using the loss function $L(\boldsymbol{\beta}, \mathbf{X}, \mathbb{P}) := \|\sum_{i=1}^n \beta_i k(x_i, \cdot) - \mathbb{E}_{x\sim\mathbb{P}}[k(x, \cdot)]\|_{\mathcal{H}}^2$. The risk of the estimator is subsequently approximated by averaging over m independent copies of $\mathbf{X}$. In this experiment, we set $n = 50$, $d = 20$, and $m = 1000$. Throughout, we use the Gaussian RBF kernel $k(x, x') = \exp(-\|x - x'\|^2/2\sigma^2)$ whose bandwidth parameter is calculated using the median heuristic, i.e., $\sigma^2 = \text{median}\{\|x_i - x_j\|^2\}$. To allow for an analytic calculation of the loss $L(\boldsymbol{\beta}, \mathbf{X}, \mathbb{P})$, we assume that the distribution $\mathbb{P}$ is a d-dimensional mixture of Gaussians [1, 8]. Specifically, the data are generated as follows: $x \sim \sum_{i=1}^4 \pi_i \mathcal{N}(\boldsymbol{\theta}_i, \Sigma_i) + \varepsilon, \theta_{ij} \sim \mathcal{U}(-10, 10), \Sigma_i \sim \mathcal{W}(3 \times \mathbf{I}_d, 7), \varepsilon \sim \mathcal{N}(0, 0.2 \times \mathbf{I}_d)$ where $\mathcal{U}(a, b)$ and $\mathcal{W}(\Sigma_0, df)$ are the uniform distribution and Wishart distribution, respectively. As in [1], we set $\boldsymbol{\pi} = [0.05, 0.3, 0.4, 0.25]$.

A natural approach for choosing λ is cross-validation procedure, which can be performed efficiently for the iterative methods such as Landweber and accelerated Landweber. For these two algorithms, we evaluate the leave-one-out score and select $\boldsymbol{\beta}^t$ at the iteration t that minimizes this score (see, e.g., Figure 3(a)). Note that these methods have the built-in property of computing the whole *regularization path* efficiently. Since each iteration of the iterated Tikhonov is in fact equivalent to the F-KMSE, we assume $t = 3$ for simplicity and use the efficient LOOCV procedure proposed in [1] to find λ at each iteration. Lastly, the truncation limit of TSVD can be identified efficiently by mean of generalized cross-validation (GCV) procedure [25]. To allow for an efficient calculation of GCV score, we resort to the alternative loss function $\mathcal{L}(\boldsymbol{\beta}) := \|\mathbf{K}\boldsymbol{\beta} - \mathbf{K}\mathbf{1}_n\|_2^2$.

Figure 3 reveals interesting aspects of the Spectral-KMSE. Firstly, as we can see in Figure 3(a), the number of iterations acts as shrinkage parameter whose optimal value can be attained within just a few iterations. Moreover, these methods do not suffer from "over-shrinking" because $\lambda \to 0$ as $t \to \infty$. In other words, if the chosen t happens to be too large, the worst we can get is the standard empirical estimator. Secondly, Figure 3(b) demonstrates that both Landweber and accelerated Landweber are more computationally efficient than the F-KMSE. Lastly, Figure 3(c) suggests that the improvement of shrinkage estimators becomes increasingly remarkable in a high-dimensional setting. Interestingly, we can observe that most Spectral-KMSE algorithms outperform the S-KMSE, which supports our hypothesis on the importance of the geometric information of RKHS mentioned in Section 3. In addition, although the TSVD still gain from shrinkage, the improvement is smaller than other algorithms. This highlights the importance of filter functions and associated parameters.

Real data. We apply Spectral-KMSE to the density estimation problem via kernel mean matching [1, 26]. The datasets were taken from the UCI repository[1] and pre-processed by standardizing each feature. Then, we fit a mixture model $Q = \sum_{j=1}^r \pi_j \mathcal{N}(\boldsymbol{\theta}_j, \sigma_j^2 \mathbf{I})$ to the pre-processed dataset

[1] http://archive.ics.uci.edu/ml/

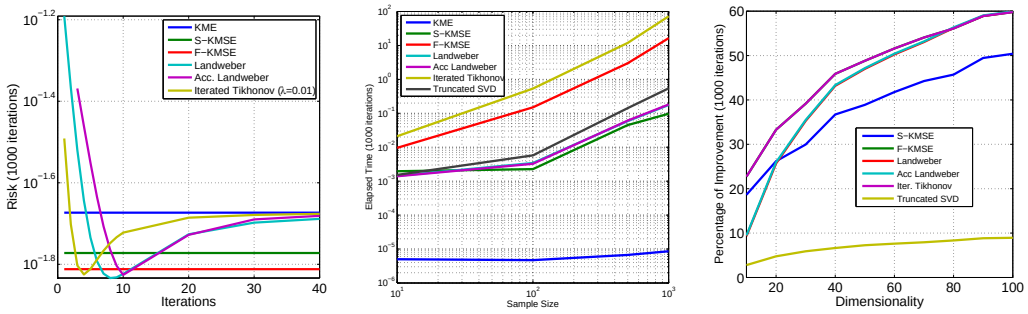

(a) risk vs. iteration (b) runtime vs. sample size (c) risk vs. dimension

Figure 3: (a) For iterative algorithms, the number of iterations acts as shrinkage parameter. (b) The iterative algorithms such as Landweber and accelerated Landweber are more efficient than the F-KMSE. (c) A percentage of improvement w.r.t. the KME, i.e., $100 \times (R - R_\lambda)/R$ where R and R_λ denote the approximated risk of KME and KMSE, respectively. Most Spectral-KMSE algorithms outperform S-KMSE which does not take into account the geometric information of the RKHS.

$\mathbf{X} := \{x_i\}_{i=1}^n$ by minimizing $\|\mu_Q - \hat{\mu}_X\|^2$ subject to the constraint $\sum_{j=1}^r \pi_j = 1$. Here μ_Q is the mean embedding of the mixture model Q and $\hat{\mu}_X$ is the empirical mean embedding obtained from $\mathbf{X}$. Based on different estimators of μ_X, we evaluate the resultant model Q by the negative log-likelihood score on the test data. The parameters $(\pi_j, \boldsymbol{\theta}_j, \sigma_j^2)$ are initialized by the best one obtained from the K-means algorithm with 50 initializations. Throughout, we set $r = 5$ and use 25% of each dataset as a test set.

Table 2: The average negative log-likelihood evaluated on the test set. The results are obtained from 30 repetitions of the experiment. The boldface represents the statistically significant results.

Dataset	KME	S-KMSE	F-KMSE	Landweber	Acc Land	Iter Tik	TSVD
ionosphere	36.1769	36.1402	36.1622	**36.1204**	**36.1554**	36.1334	36.1442
glass	10.7855	10.7403	10.7448	**10.7099**	10.7541	10.9078	10.7791
bodyfat	18.1964	18.1158	18.1810	18.1607	18.1941	18.1267	18.1061
housing	14.3016	**14.2195**	**14.0409**	14.2499	14.1983	**14.2868**	**14.3129**
vowel	13.9253	13.8426	13.8817	13.8337	14.1368	13.8633	13.8375
svmguide2	28.1091	**28.0546**	**27.9640**	28.1052	**27.9693**	**28.0417**	28.1128
vehicle	18.5295	18.3693	**18.2547**	**18.4873**	**18.3124**	**18.4128**	18.3910
wine	16.7668	**16.7548**	**16.7457**	**16.7596**	**16.6790**	16.6954	**16.5719**
wdbc	35.1916	**35.1814**	**35.0023**	**35.1402**	35.1366	35.1881	**35.1850**

Table 2 reports the results on real data. In general, the mixture model Q obtained from the proposed shrinkage estimators tend to achieve lower negative log-likelihood score than that obtained from the standard empirical estimator. Moreover, we can observe that the relative performance of different filter functions vary across datasets, suggesting that, in addition to potential gain from shrinkage, incorporating prior knowledge through the choice of filter function could lead to further improvement.

6 Conclusion

We shows that several shrinkage strategies can be adopted to improve the kernel mean estimation. This paper considers the spectral filtering approach as one of such strategies. Compared to previous work [1], our estimators take into account the specifics of kernel methods and meaningful prior knowledge through the choice of filter functions, resulting in a wider class of shrinkage estimators. The theoretical analysis also reveals a fundamental similarity to standard supervised setting. Our estimators are simple to implement and work well in practice, as evidenced by the empirical results.

Acknowledgments

The first author thanks Ingo Steinwart for pointing out existing works along the line of spectral filtering, and Arthur Gretton for suggesting the connection of shrinkage to smooth operator framework. This work was carried out when the second author was a Research Fellow in the Statistical Laboratory, Department of Pure Mathematics and Mathematical Statistics at the University of Cambridge.

References

[1] K. Muandet, K. Fukumizu, B. Sriperumbudur, A. Gretton, and B. Schölkopf. “Kernel Mean Estimation and Stein Effect”. In: *ICML*. 2014, pp. 10–18.

[2] B. Schölkopf, A. Smola, and K.-R. Müller. “Nonlinear Component Analysis as a Kernel Eigenvalue Problem”. In: *Neural Computation* 10.5 (July 1998), pp. 1299–1319.

[3] J. Shawe-Taylor and N. Cristianini. *Kernel Methods for Pattern Analysis*. Cambridge, UK: Cambridge University Press, 2004.

[4] B. Schölkopf and A. J. Smola. *Learning with Kernels: Support Vector Machines, Regularization, Optimization, and Beyond*. Cambridge, MA, USA: MIT Press, 2001.

[5] A. Berlinet and T. C. Agnan. *Reproducing Kernel Hilbert Spaces in Probability and Statistics*. Kluwer Academic Publishers, 2004.

[6] A. Smola, A. Gretton, L. Song, and B. Schölkopf. “A Hilbert Space Embedding for Distributions”. In: *ALT*. Springer-Verlag, 2007, pp. 13–31.

[7] A. Gretton, K. M. Borgwardt, M. Rasch, B. Schölkopf, and A. J. Smola. “A kernel method for the two-sample-problem”. In: *NIPS*. 2007.

[8] K. Muandet, K. Fukumizu, F. Dinuzzo, and B. Schölkopf. “Learning from Distributions via Support Measure Machines”. In: *NIPS*. 2012, pp. 10–18.

[9] L. Song, J. Huang, A. Smola, and K. Fukumizu. “Hilbert Space Embeddings of Conditional Distributions with Applications to Dynamical Systems”. In: *ICML*. 2009.

[10] K. Muandet, D. Balduzzi, and B. Schölkopf. “Domain Generalization via Invariant Feature Representation”. In: *ICML*. 2013, pp. 10–18.

[11] K. Muandet and B. Schölkopf. “One-Class Support Measure Machines for Group Anomaly Detection”. In: *UAI*. AUAI Press, 2013, pp. 449–458.

[12] B. K. Sriperumbudur, A. Gretton, K. Fukumizu, B. Schölkopf, and G. R. G. Lanckriet. “Hilbert Space Embeddings and Metrics on Probability Measures”. In: *JMLR* 99 (2010), pp. 1517–1561.

[13] K. Fukumizu, L. Song, and A. Gretton. “Kernel Bayes’ Rule: Bayesian Inference with Positive Definite Kernels”. In: *JMLR* 14 (2013), pp. 3753–3783.

[14] C. M. Stein. “Estimation of the Mean of a Multivariate Normal Distribution”. In: *The Annals of Statistics* 9.6 (1981), pp. 1135–1151.

[15] H. W. Engl, M. Hanke, and A. Neubauer. *Regularization of Inverse Problems*. Vol. 375. Mathematics and its Applications. Kluwer Academic Publishers Group, Dordrecht, 1996.

[16] E. D. Vito, L. Rosasco, and R. Verri. *Spectral Methods for Regularization in Learning Theory*. 2006.

[17] E. D. Vito, L. Rosasco, A. Caponnetto, U. D. Giovannini, and F. Odone. “Learning from Examples as an Inverse Problem.” In: *JMLR* 6 (2005), pp. 883–904.

[18] L. Baldassarre, L. Rosasco, A. Barla, and A. Verri. “Vector Field Learning via Spectral Filtering.” In: *ECML/PKDD (1)*. Vol. 6321. Lecture Notes in Computer Science. Springer, 2010, pp. 56–71.

[19] J. Kim and C. D. Scott. “Robust Kernel Density Estimation”. In: *JMLR* 13 (2012), 2529–2565.

[20] I. Steinwart and A. Christmann. *Support Vector Machines*. New York: Springer, 2008.

[21] L. Song and B. Dai. “Robust Low Rank Kernel Embeddings of Multivariate Distributions”. In: *NIPS*. 2013, pp. 3228–3236.

[22] S. Grünewälder, G. Arthur, and J. Shawe-Taylor. “Smooth Operators”. In: *ICML*. Vol. 28. 2013, pp. 1184–1192.

[23] L. L. Gerfo, L. Rosasco, F. Odone, E. D. Vito, and A. Verri. “Spectral Algorithms for Supervised Learning.” In: *Neural Computation* 20.7 (2008), pp. 1873–1897.

[24] O. V. Lepski, E. Mammen, and V. G. Spokoiny. “Optimal Spatial Adaptation to Inhomogeneous Smoothness: An Approach based on Kernel Estimates with Variable Bandwith Selectors”. In: *Annals of Statistics* 25 (1997), pp. 929–947.

[25] G. Golub, M. Heath, and G. Wahba. “Generalized Cross-Validation as a Method for Choosing a Good Ridge Parameter”. In: *Technometrics* 21 (1979), pp. 215–223.

[26] L. Song, X. Zhang, A. Smola, A. Gretton, and B. Schölkopf. “Tailoring Density Estimation via Reproducing Kernel Moment Matching”. In: *ICML*. 2008, pp. 992–999.

[27] S. Grünewälder, G. Lever, A. Gretton, L. Baldassarre, S. Patterson, and M. Pontil. “Conditional Mean Embeddings as Regressors”. In: *ICML*. 2012.

[28] L. Rosasco, M. Belkin, and E. D. Vito. “On Learning with Integral Operators”. In: *JMLR* 11 (2010), pp. 905–934.

[29] F. Bauer, S. Pereverzev, and L. Rosasco. “On Regularization Algorithms in Learning Theory”. In: *Journal of Complexity* 23.1 (2007), pp. 52–72.

[30] E. D. Vito, L. Rosasco, and A. Toigo. “Learning Sets with Separating Kernels”. In: *http://arxiv.org/abs/1204.3573* (2012).

[31] A. Gretton, K. M. Borgwardt, M. J. Rasch, B. Schölkopf, and A. Smola. “A Kernel Two-sample Test”. In: *JMLR* 13 (2012), pp. 723–773.

Appendices

A Proof of Theorem 1

(i) Since $\check{\mu}_\lambda = \hat{\mu}_{\frac{\lambda}{\lambda+1}} = \frac{\hat{\mu}_\mathbb{P}}{\lambda+1}$, we have

$$\|\check{\mu}_\lambda - \mu_\mathbb{P}\| = \left\|\frac{\hat{\mu}_\mathbb{P}}{\lambda+1} - \mu_\mathbb{P}\right\| \leq \left\|\frac{\hat{\mu}_\mathbb{P}}{\lambda+1} - \frac{\mu_\mathbb{P}}{\lambda+1}\right\| + \left\|\frac{\mu_\mathbb{P}}{\lambda+1} - \mu_\mathbb{P}\right\| \leq \|\hat{\mu}_\mathbb{P} - \mu_\mathbb{P}\| + \lambda\|\mu_\mathbb{P}\|.$$

From [7], we have that $\|\hat{\mu}_\mathbb{P} - \mu_\mathbb{P}\| = O_\mathbb{P}(n^{-1/2})$ and therefore the result follows.

(ii) Define $\Delta := \mathbb{E}_\mathbb{P}\|\hat{\mu}_\mathbb{P} - \mu_\mathbb{P}\|^2 = \frac{\int k(x,x)\,\mathrm{d}\mathbb{P}(x) - \|\mu_\mathbb{P}\|^2}{n}$. Consider

$$\begin{aligned}
\mathbb{E}_\mathbb{P}\|\check{\mu}_\lambda - \mu_\mathbb{P}\|^2 - \Delta &= \mathbb{E}_\mathbb{P}\left\|\frac{n^\beta}{n^\beta + c}(\hat{\mu}_\mathbb{P} - \mu_\mathbb{P}) - \mu_\mathbb{P}\right\|^2 - \Delta \\
&= \left(\frac{n^\beta}{n^\beta + c}\right)^2 \Delta + \frac{c^2}{(n^\beta + c)^2}\|\mu_\mathbb{P}\|^2 - \Delta \\
&= \frac{c^2\|\mu_\mathbb{P}\|^2 - (c^2 + 2cn^\beta)\Delta}{(n^\beta + c)^2}.
\end{aligned}$$

Substituting for Δ in the r.h.s. of the above equation, we have

$$\mathbb{E}_\mathbb{P}\|\check{\mu}_\lambda - \mu_\mathbb{P}\|^2 - \Delta = \frac{(nc^2 + c^2 + 2cn^\beta)\|\mu_\mathbb{P}\|^2 - (c^2 + 2cn^\beta)\int k(x,x)\,\mathrm{d}\mathbb{P}(x)}{n(n^\beta + c)^2}.$$

It is easy to verify that $\mathbb{E}_\mathbb{P}\|\check{\mu}_\lambda - \mu_\mathbb{P}\|^2 - \Delta < 0$ if

$$\frac{\|\mu_\mathbb{P}\|^2}{\int k(x,x)\,\mathrm{d}\mathbb{P}(x)} < \inf_n \frac{c^2 + 2cn^\beta}{nc^2 + c^2 + 2cn^\beta} = \frac{2^{1/\beta}\beta}{2^{1/\beta}\beta + c^{1/\beta}(\beta-1)^{(\beta-1)/\beta}}.$$

Remark. If $k(x,y) = \langle x, y\rangle$, then it is easy to check that $\mathcal{P}_{c,\beta} = \{\mathbb{P} \in M_+^1(\mathbb{R}^d) : \frac{\|\theta\|_2^2}{\mathrm{trace}(\Sigma)} < \frac{A}{1-A}\}$ where θ and Σ represent the mean vector and covariance matrix. Note that this choice of kernel yields a setting similar to classical James-Stein estimation, wherein for all n and all $\mathbb{P} \in \mathcal{P}_{c,\beta} := \{\mathbb{P} \in \mathcal{N}_{\theta,\sigma} : \|\theta\| < \sigma\sqrt{dA/(1-A)}\}$, $\check{\mu}_\lambda$ is admissible for any d, where $\mathcal{N}_{\theta,\sigma} := \{\mathbb{P} \in M_+^1(\mathbb{R}^d) : \mathrm{d}\mathbb{P}(x) = (2\pi\sigma^2)^{-d/2}e^{-\frac{\|x-\theta\|^2}{2\sigma^2}}\,\mathrm{d}x,\ \theta \in \mathbb{R}^d,\ \sigma > 0\}$. On the other hand, the James-Stein estimator is admissible for only $d \geq 3$ but for any $\mathbb{P} \in \mathcal{N}_{\theta,\sigma}$.

B Consequence of Theorem 1 if k is translation invariant

Claim: Let $k(x,y) = \psi(x-y)$, $x, y \in \mathbb{R}^d$ where ψ is a bounded continuous positive definite function with $\psi \in L^1(\mathbb{R}^d)$. For $\lambda = cn^{-\beta}$ with $c > 0$ and $\beta > 1$, define

$$\mathcal{P}_{c,\beta,\psi} := \left\{\mathbb{P} \in M_+^1(\mathbb{R}^d) : \|\phi_\mathbb{P}\|_{L^2} < \sqrt{\frac{A(2\pi)^{d/2}\psi(0)}{\|\psi\|_{L^1}}}\right\},$$

where $\phi_\mathbb{P}$ is the characteristic function of $\mathbb{P}$. Then $\forall\, n$ and $\forall\, \mathbb{P} \in \mathcal{P}_{c,\beta,\psi}$, we have $\mathbb{E}_\mathbb{P}\|\check{\mu}_\lambda - \mu_\mathbb{P}\|^2 < \mathbb{E}_\mathbb{P}\|\hat{\mu}_\mathbb{P} - \mu_\mathbb{P}\|^2$.

Proof. If $k(x,y) = \psi(x-y)$, it is easy to verify that

$$\int\int k(x,y)\,d\mathbb{P}(x)\,d\mathbb{P}(y) = \int |\phi_\mathbb{P}(\omega)|^2\widehat{\psi}(\omega)\,\mathrm{d}\omega \leq \sup_{\omega\in\mathbb{R}^d}\widehat{\psi}(\omega)\|\phi_\mathbb{P}\|_{L_2}^2 \leq (2\pi)^{-d/2}\|\psi\|_{L_1}\|\phi_\mathbb{P}\|_{L_2}^2,$$

where $\widehat{\psi}$ is the Fourier transform of ψ. On the other hand, since $|\phi_\mathbb{P}(\omega)| \leq 1$ for any $\omega \in \mathbb{R}^d$, we have

$$\begin{aligned}
\int\int k(x,y)\,d\mathbb{P}(x)\,d\mathbb{P}(y) &= \int |\phi_\mathbb{P}(\omega)|^2\widehat{\psi}(\omega)\,\mathrm{d}\omega \leq \int |\phi_\mathbb{P}(\omega)|\widehat{\psi}(\omega)\,\mathrm{d}\omega \leq \|\phi_\mathbb{P}\|_{L^2}\|\widehat{\psi}\|_{L^2} \\
&\leq \|\phi_\mathbb{P}\|_{L^2}\sqrt{\|\widehat{\psi}\|_\infty\|\widehat{\psi}\|_{L^1}} = \|\phi_\mathbb{P}\|_{L^2}\sqrt{(2\pi)^{-d/2}\|\psi\|_{L^1}\psi(0)},
\end{aligned}$$

where we used $\psi(0) = \|\widehat{\psi}\|_{L^1}$. As $\int k(x,x)\, d\mathbb{P}(x) = \psi(0)$, we have that

$$\frac{\|\mu_{\mathbb{P}}\|^2}{\int k(x,x)\, d\mathbb{P}(x)} \leq \min\left\{ \frac{\|\phi_{\mathbb{P}}\|_{L^2}^2 \|\psi\|_{L^1}}{(2\pi)^{d/2}\psi(0)}, \sqrt{\frac{\|\phi_{\mathbb{P}}\|_{L^2}^2 \|\psi\|_{L^1}}{(2\pi)^{d/2}\psi(0)}} \right\}.$$

Since $\mathbb{P} \in \mathcal{P}_{c,\beta,\psi}$, we have $\mathbb{P} \in \mathcal{P}_{c,\beta}$ and therefore the result follows. ■

C Proof of Theorem 2

Since $(e_i)_i$ is an orthonormal basis in $\mathcal{H}$, we have for any $\mathbb{P}$ and $f^* \in \mathcal{H}$

$$\mu_{\mathbb{P}} = \sum_{i=1}^{\infty} \mu_i e_i, \quad \hat{\mu}_{\mathbb{P}} = \sum_{i=1}^{\infty} \hat{\mu}_i e_i, \quad \text{and} \quad f^* = \sum_{i=1}^{\infty} f_i^* e_i,$$

where $\mu_i := \langle \mu_{\mathbb{P}}, e_i \rangle$, $\hat{\mu}_i := \langle \hat{\mu}_{\mathbb{P}}, e_i \rangle$, and $f_i^* := \langle f^*, e_i \rangle$. If follows from the Parseval's identity that

$$\begin{aligned}
\Delta &= \mathbb{E}_{\mathbb{P}} \|\hat{\mu} - \mu\|^2 = \mathbb{E}_{\mathbb{P}} \left[\sum_{i=1}^{\infty} (\hat{\mu}_i - \mu_i)^2 \right] =: \sum_{i=1}^{\infty} \Delta_i \\
\Delta_{\boldsymbol{\alpha}} &= \mathbb{E}_{\mathbb{P}} \|\hat{\mu}_{\boldsymbol{\alpha}} - \mu\|^2 = \mathbb{E}_{\mathbb{P}} \left[\sum_{i=1}^{\infty} (\alpha_i f_i^* + (1-\alpha_i)\hat{\mu}_i - \mu_i)^2 \right] =: \sum_{i=1}^{\infty} \Delta_{\boldsymbol{\alpha},i}.
\end{aligned}$$

Note that the problem has not changed and we are merely looking at it from a different perspective. To estimate $\mu_{\mathbb{P}}$, we may just as well estimate its Fourier coefficient sequence μ_i with $\hat{\mu}_i$. Based on above decomposition, we may write the risk difference $\Delta_{\boldsymbol{\alpha}} - \Delta$ as $\sum_{i=1}^{\infty} (\Delta_{\boldsymbol{\alpha},i} - \Delta_i)$. We can thus ask under which conditions on $\boldsymbol{\alpha} = (\alpha_i)$ for which $\Delta_{\boldsymbol{\alpha},i} - \Delta_i < 0$ uniformly over all i.

For each coordinate i, we have

$$\begin{aligned}
\Delta_{\boldsymbol{\alpha},i} - \Delta_i &= \mathbb{E}_{\mathbb{P}} \left[(\alpha_i f_i^* + (1-\alpha_i)\hat{\mu}_i - \mu_i)^2 \right] - \mathbb{E}_{\mathbb{P}} \left[(\hat{\mu}_i - \mu_i)^2 \right] \\
&= \mathbb{E}_{\mathbb{P}}[\alpha_i^2 f_i^2 + 2\alpha_i f_i^* (1-\alpha_i)\hat{\mu}_i + (1-\alpha_i)^2 \hat{\mu}_i^2 \\
&\quad -2\alpha_i f_i^* \mu_i - 2(1-\alpha_i)\hat{\mu}_i \mu_i + \mu_i^2] - \mathbb{E}_{\mathbb{P}}[\hat{\mu}_i^2 - 2\hat{\mu}_i \mu_i + \mu_i^2] \\
&= \alpha_i^2 f_i^2 + 2\alpha_i f_i^* \mathbb{E}_{\mathbb{P}}[\hat{\mu}_i] - 2\alpha_i^2 f_i^* \mathbb{E}_{\mathbb{P}}[\hat{\mu}_i] + (1-\alpha_i)^2 \mathbb{E}_{\mathbb{P}}[\hat{\mu}_i^2] \\
&\quad -2\alpha_i f_i^* \mu_i - 2(1-\alpha_i)\mathbb{E}_{\mathbb{P}}[\hat{\mu}_i]\mu_i + \mu_i^2 - \mathbb{E}_{\mathbb{P}}[\hat{\mu}_i^2] + 2\mu_i \mathbb{E}_{\mathbb{P}}[\hat{\mu}_i] - \mu_i^2 \\
&= \alpha_i^2 f_i^2 - 2\alpha_i^2 f_i^* \mu_i + (1-\alpha_i)^2 \mathbb{E}_{\mathbb{P}}[\hat{\mu}_i^2] - 2(1-\alpha_i)\mu_i^2 + 2\mu_i^2 - \mathbb{E}_{\mathbb{P}}[\hat{\mu}_i^2] \\
&= \alpha_i^2 f_i^2 - 2\alpha_i^2 f_i^* \mu_i + (\alpha_i^2 - 2\alpha_i)\mathbb{E}_{\mathbb{P}}[\hat{\mu}_i^2] + 2\alpha_i \mu_i^2.
\end{aligned}$$

Next, we substitute $\mathbb{E}_{\mathbb{P}}[\hat{\mu}_i^2] = \mathbb{E}_{\mathbb{P}}[(\hat{\mu}_i - \mu_i + \mu_i)^2] = \Delta_i + \mu_i^2$ into the last equation to obtain

$$\begin{aligned}
\Delta_{\boldsymbol{\alpha},i} - \Delta_i &= \alpha_i^2 f_i^2 - 2\alpha_i^2 f_i^* \mu_i + \alpha_i^2(\Delta_i + \mu_i^2) - 2\alpha_i(\Delta_i + \mu_i^2) + 2\alpha_i \mu_i^2 \\
&= \alpha_i^2 f_i^2 - 2\alpha_i^2 f_i^* \mu_i + \alpha_i^2 \Delta_i + \alpha_i^2 \mu_i^2 - 2\alpha_i \Delta_i \\
&= \alpha_i^2 (f_i^2 - 2 f_i^* \mu_i + \Delta_i + \mu_i^2) - 2\alpha_i \Delta_i \\
&= \alpha_i^2 (\Delta_i + (f_i^* - \mu_i)^2) - 2\alpha_i \Delta_i
\end{aligned}$$

which is negative if α_i satisfies

$$0 < \alpha_i < \frac{2\Delta_i}{\Delta_i + (f_i^* - \mu_i)^2}.$$

This completes the proof.

D Proof of Proposition 3

Let $\mathbf{K} = \mathbf{U}\mathbf{D}\mathbf{U}^\top$ be an eigen-decomposition of $\mathbf{K}$ where $\mathbf{U} = [\tilde{\mathbf{u}}_1, \tilde{\mathbf{u}}_2, \ldots, \tilde{\mathbf{u}}_n]$ consists of orthogonal eigenvectors of $\mathbf{K}$ such that $\mathbf{U}^\top \mathbf{U} = \mathbf{I}$ and $\mathbf{D} = \operatorname{diag}(\tilde{\gamma}_1, \tilde{\gamma}_2 \ldots, \tilde{\gamma}_n)$ consists of corresponding eigenvalues. As a result, the coefficients $\boldsymbol{\beta}(\lambda)$ can be written as

$$\boldsymbol{\beta}(\lambda) = g_\lambda(\mathbf{K})\mathbf{K}\mathbf{1}_n = \mathbf{U} g_\lambda(\mathbf{D})\mathbf{U}^\top \mathbf{K}\mathbf{1}_n = \sum_{i=1}^{n} \tilde{\mathbf{u}}_i g_\lambda(\tilde{\gamma}_i) \tilde{\mathbf{u}}_i^\top \mathbf{K}\mathbf{1}_n. \tag{9}$$

Using $\mathbf{K}\mathbf{1}_n = [\langle \hat{\mu}, k(x_1,\cdot)\rangle, \ldots, \langle \hat{\mu}, k(x_n,\cdot)\rangle]^\top$, we can rewrite (9) as

$$\begin{aligned}\boldsymbol{\beta}(\lambda) &= \sum_{i=1}^{n} \tilde{\mathbf{u}}_i g_\lambda(\tilde{\gamma}_i) \sum_{j=1}^{n} \tilde{u}_{ij} \langle \hat{\mu}, k(x_j,\cdot)\rangle \\ &= \sum_{i=1}^{n} \sqrt{\tilde{\gamma}_i}\tilde{\mathbf{u}}_i g_\lambda(\tilde{\gamma}_i) \left\langle \hat{\mu}, \frac{1}{\sqrt{\tilde{\gamma}_i}} \sum_{j=1}^{n} \tilde{u}_{ij} k(x_j,\cdot) \right\rangle,\end{aligned}$$

where $\tilde{u}_{ij}$ is the jth component of $\tilde{\mathbf{u}}_i$. Next, we invoke the relation between the eigenvectors of the matrix $\mathbf{K}$ and the eigenfunctions of the empirical covariance operator $\widehat{\mathcal{C}}_k$ in $\mathcal{H}$. That is, it is known that the ith eigenfunction of $\widehat{\mathcal{C}}_k$ can be expressed as $\tilde{\mathbf{v}}_i = (1/\sqrt{\tilde{\gamma}_i}) \sum_{j=1}^{n} \tilde{u}_{ij} k(x_j,\cdot)$ [2]. Consequently,

$$\left\langle \hat{\mu}, \frac{1}{\sqrt{\tilde{\gamma}_i}} \sum_{j=1}^{n} \tilde{u}_{ij} k(x_j,\cdot) \right\rangle = \langle \hat{\mu}, \tilde{\mathbf{v}}_i \rangle$$

and we can write the Spectral-KMSE as

$$\begin{aligned}\hat{\mu}_\lambda &= \sum_{j=1}^{n} \left[\sum_{i=1}^{n} \tilde{u}_{ij} \sqrt{\tilde{\gamma}_i} g_\lambda(\tilde{\gamma}_i) \langle \hat{\mu}, \tilde{\mathbf{v}}_i \rangle \right]_j k(x_j,\cdot) \\ &= \sum_{i=1}^{n} \sqrt{\tilde{\gamma}_i} g_\lambda(\tilde{\gamma}_i) \langle \hat{\mu}, \tilde{\mathbf{v}}_i \rangle \sum_{j=1}^{n} \tilde{u}_{ij} k(x_j,\cdot) \\ &= \sum_{i=1}^{n} g_\lambda(\tilde{\gamma}_i) \tilde{\gamma}_i \langle \hat{\mu}, \tilde{\mathbf{v}}_i \rangle \tilde{\mathbf{v}}_i.\end{aligned}$$

This completes the proof.

E Population counterpart of Spectral-KMSE

To obtain the population version of the Spectral-KMSE, we resort to the regression perspective of the kernel mean embedding which has been studied earlier in [22, 27]. The proof techniques used here are similar to those in [27]. Consider

$$\arg\min_{\mathbf{F} \in \mathcal{H}\otimes\mathcal{H}} \quad \mathbb{E}_X \left[\|k(X,\cdot) - \mathbf{F}k(X,\cdot)\|_{\mathcal{H}}^2 \right] + \lambda \|\mathbf{F}\|_{HS}^2 . \tag{10}$$

where $\mathbf{F} : \mathcal{H} \to \mathcal{H}$ is Hilbert-Schmidt. We can expand the regularized loss (10) as

$$\begin{aligned}&\mathbb{E}_X \left[\|k(X,\cdot) - \mathbf{F}k(X,\cdot)\|_{\mathcal{H}}^2 \right] + \lambda \|\mathbf{F}\|_{HS}^2 \\ &= \mathbb{E}_X \langle k(X,\cdot), k(X,\cdot)\rangle_{\mathcal{H}} - 2\mathbb{E}_X \langle k(X,\cdot), \mathbf{F}k(X,\cdot)\rangle_{\mathcal{H}} + \mathbb{E}_X \langle \mathbf{F}k(X,\cdot), \mathbf{F}k(X,\cdot)\rangle_{\mathcal{H}} + \lambda \langle \mathbf{F}, \mathbf{F}\rangle_{HS} \\ &= \mathbb{E}_X \langle k(X,\cdot), k(X,\cdot)\rangle_{\mathcal{H}} - 2\mathbb{E}_X \langle k(X,\cdot) \otimes k(X,\cdot), \mathbf{F}\rangle_{HS} + \mathbb{E}_X \langle k(X,\cdot), \mathbf{F}^*\mathbf{F}k(X,\cdot)\rangle_{\mathcal{H}} + \lambda \langle \mathbf{F}, \mathbf{F}\rangle_{HS} \\ &= \mathbb{E}_X \langle k(X,\cdot), k(X,\cdot)\rangle_{\mathcal{H}} - 2\langle \mathcal{C}_k, \mathbf{F}\rangle_{HS} + \langle \mathcal{C}_k, \mathbf{F}^*\mathbf{F}\rangle_{HS} + \lambda \langle \mathbf{F}, \mathbf{F}\rangle_{HS},\end{aligned}$$

where $\mathbf{F}^*$ denotes the adjoint of $\mathbf{F}$ and $\mathcal{C}_k = \mathbb{E}_X[k(X,\cdot) \otimes k(X,\cdot)]$. Next, we show that the solution to the above expression is $\mathbf{F} := \mathcal{C}_k(\mathcal{C}_k + \lambda \mathbf{I})^{-1}$. Defining $\mathbf{A} := \mathbf{F}(\mathcal{C}_k + \lambda \mathbf{I})^{1/2}$, the above expression can be rewritten as

$$\begin{aligned}&\mathbb{E}_X \langle k(X,\cdot), k(X,\cdot)\rangle_{\mathcal{H}} - 2\langle \mathcal{C}_k, \mathbf{F}\rangle_{HS} + \langle \mathcal{C}_k, \mathbf{F}^*\mathbf{F}\rangle_{HS} + \lambda \langle \mathbf{F}, \mathbf{F}\rangle_{HS} \\ &= \mathbb{E}_X \langle k(X,\cdot), k(X,\cdot)\rangle_{\mathcal{H}} - 2\langle \mathcal{C}_k, \mathbf{F}\rangle_{HS} + \langle \mathcal{C}_k + \lambda \mathbf{I}, \mathbf{F}^*\mathbf{F}\rangle_{HS} \\ &= \mathbb{E}_X \langle k(X,\cdot), k(X,\cdot)\rangle_{\mathcal{H}} - 2\langle \mathcal{C}_k, \mathbf{F}\rangle_{HS} + \left\langle \mathbf{F}(\mathcal{C}_k + \lambda \mathbf{I})^{1/2}, \mathbf{F}(\mathcal{C}_k + \lambda \mathbf{I})^{1/2} \right\rangle_{HS} \\ &= \mathbb{E}_X \langle k(X,\cdot), k(X,\cdot)\rangle_{\mathcal{H}} - 2\langle \mathcal{C}_k, \mathbf{A}(\mathcal{C}_k + \lambda \mathbf{I})^{-1/2}\rangle_{HS} + \langle \mathbf{A}, \mathbf{A}\rangle_{HS} \\ &= \mathbb{E}_X \langle k(X,\cdot), k(X,\cdot)\rangle_{\mathcal{H}} - \left\| \mathcal{C}_k(\mathcal{C}_k + \lambda \mathbf{I})^{-1/2} \right\|_{HS}^2 + \left\| \mathcal{C}_k(\mathcal{C}_k + \lambda \mathbf{I})^{-1/2} - A \right\|_{HS}^2 .\end{aligned}$$

As a result, the above expression is minimized when $\mathbf{A} = \mathcal{C}_k(\mathcal{C}_k + \lambda \mathbf{I})^{-1/2}$, implying that $\mathbf{F} = \mathcal{C}_k(\mathcal{C}_k + \lambda \mathbf{I})^{-1}$. As in the sample case, a natural estimate of the Spectral-KMSE is

$$\mu_\lambda = \mathbf{F}\mu_{\mathbb{P}} = \mathcal{C}_k(\mathcal{C}_k + \lambda \mathbf{I})^{-1}\mu_{\mathbb{P}}.$$

F Proof of Proposition 4

The proof employs the relation between the Gram matrix $\mathbf{K}$ and the empirical covariance operator $\widehat{\mathcal{C}}_k$ shown in Lemma 3. It is known that the operator $\widehat{\mathcal{C}}_k$ is of finite rank, self-adjoint, and positive. Moreover, its spectrum has only finitely many nonzero elements [28]. If $\tilde{\gamma}_i$ is a nonzero eigenvalue and $\tilde{\mathbf{v}}_i$ is the corresponding eigenfunction of $\widehat{\mathcal{C}}_k$, then the following decomposition holds

$$\widehat{\mathcal{C}}_k f = \sum_{i=1}^{n} \tilde{\gamma}_i \langle f, \tilde{\mathbf{v}}_i \rangle_{\mathcal{H}} \tilde{\mathbf{v}}_i, \quad \forall f \in \mathcal{H}.$$

Note that it may be that $k < n$ where k is the rank of $\widehat{\mathcal{C}}_k$. In that case, the above decomposition still holds. Setting $f = \hat{\mu}$ and applying the definition of the filter function g_λ to the operator $\widehat{\mathcal{C}}_k$ yield

$$\hat{\mu}_\lambda = \widehat{\mathcal{C}}_k g_\lambda(\widehat{\mathcal{C}}_k)\hat{\mu} = \sum_{i=1}^{n} g_\lambda(\tilde{\gamma}_i)\tilde{\gamma}_i \langle \hat{\mu}, \tilde{\mathbf{v}}_i \rangle_{\mathcal{H}} \tilde{\mathbf{v}}_i,$$

which is exactly the decomposition given in Lemma 3. This completes the proof.

G Proof of Theorem 5

Consider the following decomposition

$$\begin{aligned}
\hat{\mu}_\lambda - \mu_{\mathbb{P}} &= \widehat{\mathcal{C}}_k g_\lambda(\widehat{\mathcal{C}}_k)\hat{\mu}_{\mathbb{P}} - \mu_{\mathbb{P}} \\
&= \widehat{\mathcal{C}}_k g_\lambda(\widehat{\mathcal{C}}_k)(\hat{\mu}_{\mathbb{P}} - \mu_{\mathbb{P}}) + \widehat{\mathcal{C}}_k g_\lambda(\widehat{\mathcal{C}}_k)\mu_{\mathbb{P}} - \mu_{\mathbb{P}} \\
&= \widehat{\mathcal{C}}_k g_\lambda(\widehat{\mathcal{C}}_k)(\hat{\mu}_{\mathbb{P}} - \mu_{\mathbb{P}}) + (\widehat{\mathcal{C}}_k g_\lambda(\widehat{\mathcal{C}}_k) - I)\widehat{\mathcal{C}}_k^{\beta} h + (\widehat{\mathcal{C}}_k g_\lambda(\widehat{\mathcal{C}}_k) - I)(\mathcal{C}_k^{\beta} - \widehat{\mathcal{C}}_k^{\beta})h
\end{aligned}$$

where we used the fact that there exists $h \in \mathcal{H}$ such that $\mu_{\mathbb{P}} = \mathcal{C}_k^{\beta} h$ as we assumed that $\mu_{\mathbb{P}} \in \mathcal{R}(\mathcal{C}_k^{\beta})$ for some $\beta > 0$. Therefore

$$\|\hat{\mu}_\lambda - \mu_{\mathbb{P}}\| \leq \|\widehat{\mathcal{C}}_k g_\lambda(\widehat{\mathcal{C}}_k)\|_{op}\|\hat{\mu}_{\mathbb{P}} - \mu_{\mathbb{P}}\| + \|(\widehat{\mathcal{C}}_k g_\lambda(\widehat{\mathcal{C}}_k) - I)\widehat{\mathcal{C}}_k^{\beta}\|_{op}\|h\| + \|\widehat{\mathcal{C}}_k g_\lambda(\widehat{\mathcal{C}}_k) - I\|_{op}\|\mathcal{C}_k^{\beta} - \widehat{\mathcal{C}}_k^{\beta}\|_{op}\|h\|$$

where we used the fact that $\|Ab\| \leq \|A\|_{op}\|b\|$ with $A : \mathcal{H} \to \mathcal{H}$ being a bounded operator, $b \in \mathcal{H}$ and $\|\cdot\|_{op}$ denoting the operator norm defined as $\|A\|_{op} := \sup\{\|Ab\| : \|b\| = 1\}$.

By $(C1)$, $(C2)$ and $(C3)$, we have $\|\widehat{\mathcal{C}}_k g_\lambda(\widehat{\mathcal{C}}_k)\|_{op} \leq B$, $\|\widehat{\mathcal{C}}_k g_\lambda(\widehat{\mathcal{C}}_k) - I\|_{op} \leq C$ and $\|(\widehat{\mathcal{C}}_k g_\lambda(\widehat{\mathcal{C}}_k) - I)\widehat{\mathcal{C}}_k^{\beta}\|_{op} \leq D\lambda^{\min\{\beta,\eta_0\}}$ respectively. Denoting $\|h\| = \|\mathcal{C}_k^{-\beta}\mu_{\mathbb{P}}\|$, we therefore have

$$\|\hat{\mu}_\lambda - \mu_{\mathbb{P}}\| \leq B\|\hat{\mu}_{\mathbb{P}} - \mu_{\mathbb{P}}\| + D\lambda^{\min\{\beta,\eta_0\}}\|\mathcal{C}_k^{-\beta}\mu_{\mathbb{P}}\| + C\|\mathcal{C}_k^{\beta} - \widehat{\mathcal{C}}_k^{\beta}\|_{op}\|\mathcal{C}_k^{-\beta}\mu_{\mathbb{P}}\|. \tag{11}$$

For $0 \leq \beta \leq 1$, it follows from Theorem 1 in [29] that there exists a constant τ_1 such that

$$\|\mathcal{C}_k^{\beta} - \widehat{\mathcal{C}}_k^{\beta}\|_{op} \leq \tau_1\|\mathcal{C}_k - \widehat{\mathcal{C}}_k\|_{op}^{\beta} \leq \tau_1\|\mathcal{C}_k - \widehat{\mathcal{C}}_k\|_{HS}^{\beta}.$$

On the other hand, since $\alpha \mapsto \alpha^{\beta}$ is Lipschitz on $[0, \kappa^2]$ for $\beta \geq 1$, the following lemma yields that

$$\|\mathcal{C}_k^{\beta} - \widehat{\mathcal{C}}_k^{\beta}\|_{op} \leq \|\mathcal{C}_k^{\beta} - \widehat{\mathcal{C}}_k^{\beta}\|_{HS} \leq \tau_2\|\mathcal{C}_k - \widehat{\mathcal{C}}_k\|_{HS}$$

where τ_2 is the Lipschitz constant of $\alpha \mapsto \alpha^{\beta}$ on $[0, \kappa^2]$. In other words,

$$\|\mathcal{C}_k^{\beta} - \widehat{\mathcal{C}}_k^{\beta}\|_{op} \leq \max\{\tau_1, \tau_2\}\|\mathcal{C}_k - \widehat{\mathcal{C}}_k\|_{HS}^{\min\{1,\beta\}}. \tag{12}$$

Lemma 6 (Contributed by Anreas Maurer, see Lemma 5 in [30]). *Suppose A and B are self-adjoint Hilbert-Schmidt operators on a separable Hilbert space H with spectrum contained in the interval $[a, b]$, and let $(\sigma_i)_{i\in I}$ and $(\tau_j)_{j\in J}$ be the eigenvalues of A and B, respectively. Given a function $r : [a, b] \to \mathbb{R}$, if there exists a finite constant L such that*

$$|r(\sigma_i) - r(\tau_j)| \leq L|\sigma_i - \tau_j|, \ \forall\, i \in I,\ j \in J,$$

then

$$\|r(A) - r(B)\|_{HS} \leq L\|A - B\|_{HS}.$$

Using (12) in (11), we have

$$\|\hat{\mu}_\lambda - \mu_{\mathbb{P}}\| \leq B\|\hat{\mu}_{\mathbb{P}} - \mu_{\mathbb{P}}\| + D\lambda^{\min\{\beta,\eta_0\}}\|\mathcal{C}_k^{-\beta}\mu_{\mathbb{P}}\| + C\tau\|\mathcal{C}_k - \widehat{\mathcal{C}}_k\|_{HS}^{\min\{1,\beta\}}\|\mathcal{C}_k^{-\beta}\mu_{\mathbb{P}}\|, \quad (13)$$

where $\tau := \max\{\tau_1, \tau_2\}$. We now obtain bounds on $\|\hat{\mu}_{\mathbb{P}} - \mu_{\mathbb{P}}\|$ and $\|\mathcal{C}_k - \widehat{\mathcal{C}}_k\|_{HS}$ using the following results.

Lemma 7 ([31]). *Suppose that $\kappa = \sup_{x\in\mathcal{X}} \sqrt{k(x,x)}$. For any $\delta > 0$, the following inequality holds with probability at least $1 - e^{-\delta}$*

$$\|\hat{\mu}_{\mathbb{P}} - \mu_{\mathbb{P}}\| \leq \frac{2\kappa + \kappa\sqrt{2\delta}}{\sqrt{n}}.$$

Lemma 8 (e.g., see Theorem 7 in [28]). *Let $\kappa := \sup_{x\in\mathcal{X}} \sqrt{k(x,x)}$. For $n \in \mathbb{N}$ and any $\delta > 0$, the following inequality holds with probability at least $1 - 2e^{-\delta}$:*

$$\left\|\widehat{\mathcal{C}}_k - \mathcal{C}_k\right\|_{HS} \leq \frac{2\sqrt{2}\kappa^2\sqrt{\delta}}{\sqrt{n}}.$$

Using Lemmas 7 and 8 in (13), for any $\delta > 0$, with probability $1 - 3e^{-\delta}$, we obtain

$$\|\hat{\mu}_\lambda - \mu_{\mathbb{P}}\| \leq \frac{2\kappa B + \kappa B\sqrt{2\delta}}{\sqrt{n}} + D\lambda^{\min\{\beta,\eta_0\}}\|\mathcal{C}_k^{-\beta}\mu_{\mathbb{P}}\| + C\tau\frac{(2\sqrt{2}\kappa^2\sqrt{\delta})^{\min\{1,\beta\}}}{n^{\min\{1/2,\beta/2\}}}\|\mathcal{C}_k^{-\beta}\mu_{\mathbb{P}}\|.$$

H Shrinkage parameter $\lambda = cn^{-\beta}$

In this section, we provide supplementary results that demonstrate the effect of the shrinkage parameter λ presented in Theorem 1. That is, if we choose $\lambda = cn^{-\beta}$ for some $c > 0$ and $\beta > 1$, the estimator $\check{\mu}_\lambda$ is a proper estimator of μ. Unfortunately, the true value of β, which characterizes the smoothness of the true kernel mean $\mu_{\mathbb{P}}$, is not known in practice. Nevertheless, we provide simulated experiments that illustrate the convergence of the estimator $\check{\mu}_\lambda$ for different values of c and β.

The data-generating distribution used in this experiment is identical to the one we consider in our previous experiments on synthetic data. That is, the data are generated as follows: $x \sim \sum_{i=1}^4 \pi_i\mathcal{N}(\boldsymbol{\theta}_i, \Sigma_i) + \varepsilon, \theta_{ij} \sim \mathcal{U}(-10, 10), \Sigma_i \sim \mathcal{W}(3 \times \mathbf{I}_d, 7), \varepsilon \sim \mathcal{N}(0, 0.2 \times \mathbf{I}_d)$ where $\mathcal{U}(a, b)$ and $\mathcal{W}(\Sigma_0, df)$ are the uniform distribution and Wishart distribution, respectively. We set $\boldsymbol{\pi} = [0.05, 0.3, 0.4, 0.25]$. We use the Gaussian RBF kernel $k(x, x') = \exp(-\|x-x'\|^2/2\sigma^2)$ whose bandwidth parameter is calculated using the median heuristic, i.e., $\sigma^2 = \text{median}\{\|x_i - x_j\|^2\}$. Figure 4 depicts the comparisons between the standard kernel mean estimator and the shrinkage estimators with varying values of c and β.

As we can see in Figure 4, if c is very small or β is very large, the shrinkage estimator $\check{\mu}_\lambda$ behaves like the empirical estimator $\hat{\mu}_{\mathbb{P}}$. This coincides with the intuition given in Theorem 1. Note that the value of β specifies the smoothness of the true kernel mean μ and is unknown in practice. Thus, one of the interesting future directions is to develop procedure that can adapt to this unknown parameter automatically.

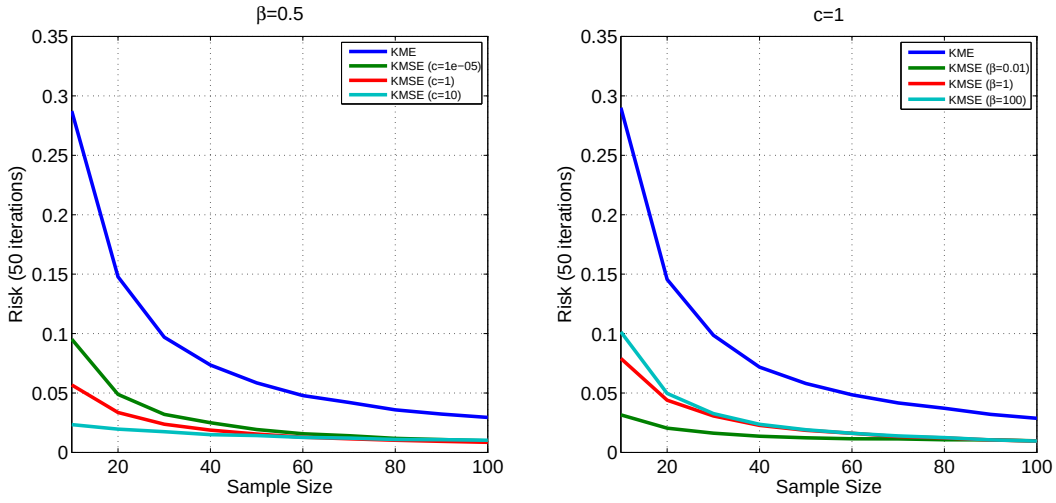

Figure 4: The risk of shrinkage estimator $\check{\mu}_\lambda$ when $\lambda = cn^{-\beta}$. The left figure shows the risk of the shrinkage estimator as sample size increases while fixing the value of β, whereas the right figure shows the same plots while fixing the value of c. See text for more explanation.